

Experimenting with Transitive Verbs in a DisCoCat

Edward Grefenstette

University of Oxford

Department of Computer Science

Wolfson Building, Parks Road

Oxford OX1 3QD, UK

edward.grefenstette@cs.ox.ac.uk

Mehrnoosh Sadrzadeh

University of Oxford

Department of Computer Science

Wolfson Building, Parks Road

Oxford OX1 3QD, UK

mehrs@cs.ox.ac.uk

Abstract

Formal and distributional semantic models offer complementary benefits in modeling meaning. The categorical compositional distributional model of meaning of Coecke et al. (2010) (abbreviated to DisCoCat in the title) combines aspects of both to provide a general framework in which meanings of words, obtained distributionally, are composed using methods from the logical setting to form sentence meaning. Concrete consequences of this general abstract setting and applications to empirical data are under active study (Grefenstette et al., 2011; Grefenstette and Sadrzadeh, 2011). In this paper, we extend this study by examining transitive verbs, represented as matrices in a DisCoCat. We discuss three ways of constructing such matrices, and evaluate each method in a disambiguation task developed by Grefenstette and Sadrzadeh (2011).

are active topics in current natural language semantics research, e.g. see (Grefenstette et al., 2011; Grefenstette and Sadrzadeh, 2011; Clark et al., 2010; Baroni and Zamparelli, 2010; Guevara, 2010; Mitchell and Lapata, 2008).

In this paper, we focus on our recent concrete DisCoCat model (Grefenstette and Sadrzadeh, 2011) and in particular on nouns composed with transitive verbs. Whereby the meaning of a transitive sentence ‘sub tverb obj’ is obtained by taking the component-wise multiplication of the matrix of the verb with the kronecker product of the vectors of subject and object:

$$\overrightarrow{\text{sub tverb obj}} = \underline{\text{tverb}} \odot (\overrightarrow{\text{sub}} \otimes \overrightarrow{\text{obj}}) \quad (1)$$

In most logical models, transitive verbs are modeled as relations; in the categorical model the relational nature of such verbs gets manifested in their matrix representation: if subject and object are each r -dimensional row vectors in some space N , the verb will be a $r \times r$ matrix in the space $N \otimes N$. There are different ways of learning the weights of this matrix. In (Grefenstette and Sadrzadeh, 2011), we developed and implemented one such method on the data from the British National Corpus. The matrix of each verb was constructed by taking the sum of the kronecker products of all of the subject/object pairs linked to that verb in the corpus. We refer to this method as the *indirect method*. This is because the weight c_{ij} is obtained from the weights of the subject and object vectors (computed via co-occurrence with bases \vec{n}_i and \vec{n}_j respectively), rather than directly from the context of the verb itself, as would be the case in lexical distributional models. This

1 Background

The categorical distributional compositional model of meaning of Coecke et al. (2010) combines the modularity of formal semantic models with the empirical nature of vector space models of lexical semantics. The meaning of a sentence is defined to be the application of its grammatical structure—represented in a type-logical model—to the kronecker product of the meanings of its words, as computed in a distributional model. The concrete and experimental consequences of this setting, and other models that aim to bring together the logical and distributional approaches,

construction method was evaluated against an extension of Mitchell and Lapata (2008)’s disambiguation task from intransitive to transitive sentences. We showed and discussed how and why our method, which is moreover scalable and respects the grammatical structure of the sentence, resulted in better results than other known models of semantic vector composition.

As a motivation for the current paper, note that there are at least two different factors at work in Equation (1): one is the matrix of the verb, denoted by $\underline{\text{tverb}}$, and the other is the kronecker product of subject and object vectors $\vec{\text{sub}} \otimes \vec{\text{obj}}$. Our model’s mathematical formulation of composition prohibits us from changing the latter kronecker product, but the ‘content’ of the verb matrices can be built through different procedures.

In recent work we used a standard lexical distributional model for nouns and engineered our verbs to have a more sophisticated structure because of the higher dimensional space they occupy. In particular, we argued that the resulting matrix of the verb should represent ‘the extent according to which the verb has related the properties of subjects to the properties of its objects’, developed a general procedure to build such matrices, then studied their empirical consequences. One question remained open: what would be the consequence of starting from the standard lexical vector of the verb, then encoding it into the higher dimensional space using different (possibly ad-hoc but nonetheless interesting) mathematically inspired methods.

In a nutshell, the lexical vector of the verb is denoted by $\vec{\text{tverb}}$ and similar to vectors of subject and object, it is an r -dimensional row vector. Since the kronecker product of subject and object ($\vec{\text{sub}} \otimes \vec{\text{obj}}$) is $r \times r$, in order to make $\vec{\text{tverb}}$ applicable in Equation 1, i.e. to be able to substitute it for $\underline{\text{tverb}}$, we need to encode it into a $r \times r$ matrix in the $N \otimes N$ space. In what follows, we investigate the empirical consequences of three different encodings methods.

2 From Vectors to Matrices

In this section, we discuss three different ways of encoding r dimensional lexical verb vectors into $r \times r$ verb matrices, and present empirical results for each. We use the additional structure that the kronecker

product provides to represent the relational nature of transitive verbs. The results are an indication that the extra information contained in this larger space contributes to higher quality composition.

One way to encode an r -dimensional vector as a $r \times r$ matrix is to embed it as the diagonal of that matrix. It remains open to decide what the non-diagonal values should be. We experimented with 0s and 1s as padding values. If the vector of the verb is $[c_1, c_2, \dots, c_r]$ then for the 0 case (referred to as **0-diag**) we obtain the following matrix:

$$\underline{\text{tverb}} = \begin{pmatrix} c_1 & 0 & \dots & 0 \\ 0 & c_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & c_r \end{pmatrix}$$

For the 1 case (referred to as **1-diag**) we obtain the following matrix:

$$\underline{\text{tverb}} = \begin{pmatrix} c_1 & 1 & \dots & 1 \\ 1 & c_2 & \dots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \dots & c_r \end{pmatrix}$$

We also considered a third case where the vector is encoded into a matrix by taking the kronecker product of the verb vector with itself:

$$\underline{\text{tverb}} = \vec{\text{tverb}} \otimes \vec{\text{tverb}}$$

So for $\vec{\text{tverb}} = [c_1, c_2, \dots, c_r]$ we obtain the following matrix:

$$\underline{\text{tverb}} = \begin{pmatrix} c_1 c_1 & c_1 c_2 & \dots & c_1 c_r \\ c_2 c_1 & c_2 c_2 & \dots & c_2 c_r \\ \vdots & \vdots & \ddots & \vdots \\ c_r c_1 & c_r c_2 & \dots & c_r c_r \end{pmatrix}$$

3 Degrees of synonymity for sentences

The degree of synonymity between meanings of two sentences is computed by measuring their geometric distance. In this work, we used the cosine measure. For two sentences ‘sub₁ tverb₁ obj₁’ and ‘sub₂ tverb₂ obj₂’, this is obtained by taking the *Frobenius* inner product of $\vec{\text{sub}}_1 \vec{\text{tverb}}_1 \vec{\text{obj}}_1$ and $\vec{\text{sub}}_2 \vec{\text{tverb}}_2 \vec{\text{obj}}_2$. The use of *Frobenius* product rather than the dot product is because the calculation in Equation (1) produces matrices rather than

row vectors. We normalized the outputs by the multiplication of the lengths of their corresponding matrices.

4 Experiment

In this section, we describe the experiment used to evaluate and compare these three methods. The experiment is on the dataset developed in (Grefenstette and Sadrzadeh, 2011).

Parameters We used the parameters described by Mitchell and Lapata (2008) for the noun and verb vectors. All vectors were built from a lemmatised version of the BNC. The noun basis was the 2000 most common context words, basis weights were the probability of context words given the target word divided by the overall probability of the context word. These features were chosen to enable easy comparison of our experimental results with those of Mitchell and Lapata’s original experiment, in spite of the fact that there may be more sophisticated lexical distributional models available.

Task This is an extension of Mitchell and Lapata (2008)’s disambiguation task from intransitive to transitive sentences. The general idea behind the transitive case (similar to the intransitive one) is as follows: meanings of ambiguous transitive verbs vary based on their subject-object context. For instance the verb ‘meet’ means ‘satisfied’ in the context ‘the system met the criterion’ and it means ‘visit’, in the context ‘the child met the house’. Hence if we build meaning vectors for these sentences compositionally, the degrees of synonymy of the sentences can be used to disambiguate the meanings of the verbs in them.

Suppose a verb has two meanings a and b and that it has occurred in two sentences. Then if in both of these sentences it has its meaning a , the two sentences will have a high degree of synonymy, whereas if in one sentence the verb has meaning a and in the other meaning b , the sentences will have a lower degree of synonymy. For instance ‘the system met the criterion’ and ‘the system satisfied the criterion’ have a high degree of semantic similarity, and similarly for ‘the child met the house’ and ‘the child visited the house’. This degree decreases for the pair ‘the child met the house’ and ‘the child sat-

isfied the house’.

Dataset The dataset is built using the same guidelines as Mitchell and Lapata (2008), using transitive verbs obtained from CELEX¹ paired with subjects and objects. We first picked 10 transitive verbs from the most frequent verbs of the BNC. For each verb, two different non-overlapping meanings were retrieved, by using the JCN (Jiang Conrath) information content synonymy measure of WordNet to select maximally different synsets. For instance for ‘meet’ we obtained ‘visit’ and ‘satisfy’. For each original verb, ten sentences containing that verb with the same role were retrieved from the BNC. Examples of such sentences are ‘the system met the criterion’ and ‘the child met the house’. For each such sentence, we generated two other related sentences by substituting their verbs by each of their two synonyms. For instance, we obtained ‘the system satisfied the criterion’ and ‘the system visited the criterion’ for the first meaning and ‘the child satisfied the house’ and ‘the child visited the house’ for the second meaning. This procedure provided us with 200 pairs of sentences.

The dataset was split into four non-identical sections of 100 entries such that each sentence appears in exactly two sections. Each section was given to a group of evaluators who were asked to assign a similarity score to simple transitive sentence pairs formed from the verb, subject, and object provided in each entry (*e.g.* ‘the system met the criterion’ from ‘system meet criterion’). The scoring scale for human judgement was [1, 7], where 1 was most dissimilar and 7 most identical.

Separately from the group annotation, each pair in the dataset was given the additional arbitrary classification of HIGH or LOW similarity by the authors.

Evaluation Method To evaluate our methods, we first applied our formulae to compute the similarity of each phrase pair on a scale of [0, 1] and then compared it with human judgement of the same pair. The comparison was performed by measuring Spearman’s ρ , a rank correlation coefficient ranging from -1 to 1. This provided us with the degree of correlation between the similarities as computed by our model and as judged by human evaluators.

¹<http://celex.mpi.nl/>

Following Mitchell and Lapata (2008), we also computed the mean of HIGH and LOW scores. However, these scores were solely based on the authors’ personal judgements and as such (and on their own) do not provide a very reliable measure. Therefore, like Mitchell and Lapata (2008), the models were ultimately judged by Spearman’s ρ .

The results are presented in table 4. The additive and multiplicative rows have, as composition operation, vector addition and component-wise multiplication. The *Baseline* is from a non-compositional approach; it is obtained by comparing the verb vectors of each pair directly and ignoring their subjects and objects. The *UpperBound* is set to be inter-annotator agreement.

Model	High	Low	ρ
Baseline	0.47	0.44	0.16
Add	0.90	0.90	0.05
Multiply	0.67	0.59	0.17
Categorical			
Indirect matrix	0.73	0.72	0.21
0-diag matrix	0.67	0.59	0.17
1-diag matrix	0.86	0.85	0.08
$v \otimes v$ matrix	0.34	0.26	0.28
UpperBound	4.80	2.49	0.62

Table 1: Results of compositional disambiguation.

The **indirect matrix** performed better than the vectors encoded in diagonal matrices padded with 0 and 1. However, surprisingly, the kronecker product of this vector with itself provided better results than all the above. The results were statistically significant with $p < 0.05$.

5 Analysis of the Results

Suppose the vector of subject is $[s_1, s_2, \dots, s_r]$ and the vector of object is $\vec{obj} = [o_1, o_2, \dots, o_r]$, then the matrix of $\vec{sub} \otimes \vec{obj}$ is:

$$\begin{pmatrix} s_1 o_1 & s_1 o_2 & \cdots & s_1 o_r \\ s_2 o_1 & s_2 o_2 & \cdots & s_2 o_r \\ \vdots & \vdots & \ddots & \vdots \\ s_r o_1 & s_r o_2 & \cdots & s_r o_r \end{pmatrix}$$

After computing Equation (1) for each generation method of tverb, we obtain the following three ma-

trices for the meaning of a transitive sentence:

$$\mathbf{0-diag}: \begin{pmatrix} c_1 s_1 o_1 & 0 & \cdots & 0 \\ 0 & c_2 s_2 o_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & c_r s_r o_r \end{pmatrix}$$

This method discards all of the non-diagonal information about the subject and object, for example there is no occurrence of $s_1 o_2$, $s_2 o_1$, etc.

$$\mathbf{1-diag}: \begin{pmatrix} c_1 s_1 o_1 & s_1 o_2 & \cdots & s_1 o_r \\ s_2 o_1 & c_2 s_2 o_2 & \cdots & s_2 o_r \\ \vdots & \vdots & \ddots & \vdots \\ s_r o_1 & s_r o_2 & \cdots & c_r s_r o_r \end{pmatrix}$$

This method conserves the information about the subject and object, but only applies the information of the verb to the diagonals: s_1 and o_2 , s_2 and o_1 , etc. are never related to each other via the verb.

$$v \otimes v: \begin{pmatrix} c_1 c_1 s_1 o_1 & c_1 c_2 s_1 o_2 & \cdots & c_1 c_r s_1 o_r \\ c_2 c_1 s_2 o_1 & c_2 c_2 s_2 o_2 & \cdots & c_2 c_r s_2 o_r \\ \vdots & \vdots & \ddots & \vdots \\ c_r c_1 s_r o_1 & c_r c_2 s_r o_2 & \cdots & c_r c_r s_r o_r \end{pmatrix}$$

This method not only conserves the information of the subject and object, but also applies to them all of the information encoded in the verb. These data propagate to *Frobenius* products when computing the semantic similarity of sentences and justify the empirical results.

The unexpectedly good performance of the $v \otimes v$ matrix relative to the more complex indirect method is surprising, and certainly demands further investigation. What is sure is that they each draw upon different aspects of semantic composition to provide better results. There is certainly room for improvement and empirical optimisation in both of these relation-matrix construction methods.

Furthermore, the success of both of these methods relative to the others examined in Table 1 shows that it is the extra information provided in the matrix (rather than just the diagonal, representing the lexical vector) that encodes the *relational nature* of transitive verbs, thereby validating in part the requirement suggested in Coecke et al. (2010) and Grefenstette and Sadrzadeh (2011) that relational word vectors live in a space the dimensionality of which be a function of the arity of the relation.

References

- [Alshawil992] H. Alshawil (ed). 1992. *The Core Language Engine*. MIT Press.
- [Baroni and Zamparelli2010] M. Baroni and R. Zamparelli. 2010. *Nouns are vectors, adjectives are matrices*. Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP).
- [Clark et al.2010] D. Clarke, R. Lutz and D. Weir. 2010. *Semantic Composition with Quotient Algebras*. Proceedings of Geometric Models of Natural Language Semantics (GEMS-2010).
- [Clark and Pulman2007] S. Clark and S. Pulman. 2007. *Combining Symbolic and Distributional Models of Meaning*. Proceedings of AAAI Spring Symposium on Quantum Interaction. AAAI Press.
- [Coecke et al.2010] B. Coecke, M. Sadrzadeh and S. Clark. 2010. *Mathematical Foundations for Distributed Compositional Model of Meaning*. Lambek Festschrift. Linguistic Analysis **36**, 345–384. J. van Benthem, M. Moortgat and W. Buszkowski (eds.).
- [Curran2004] J. Curran. 2004. *From Distributional to Semantic Similarity*. PhD Thesis, University of Edinburgh.
- [Erk and Pado2008] K. Erk and S. Padó. 2004. *A Structured Vector Space Model for Word Meaning in Context*. Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP), 897–906.
- [Frege1892] G. Frege 1892. *Über Sinn und Bedeutung*. Zeitschrift für Philosophie und philosophische Kritik 100.
- [Firth1957] J. R. Firth. 1957. *A synopsis of linguistic theory 1930-1955*. Studies in Linguistic Analysis.
- [Grefenstette et al.2011] E. Grefenstette, M. Sadrzadeh, S. Clark, B. Coecke, S. Pulman. 2011. *Concrete Compositional Sentence Spaces for a Compositional Distributional Model of Meaning*. International Conference on Computational Semantics (IWCS'11). Oxford.
- [Grefenstette and Sadrzadeh2011] E. Grefenstette, M. Sadrzadeh. 2011. *Experimental Support for a Categorical Compositional Distributional Model of Meaning*. Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing.
- [Grefenstette1994] G. Grefenstette. 1994. *Explorations in Automatic Thesaurus Discovery*. Kluwer.
- [Guevara2010] E. Guevara. 2010. *A Regression Model of Adjective-Noun Compositionality in Distributional Semantics*. Proceedings of the ACL GEMS Workshop.
- [Harris1966] Z. S. Harris. 1966. *A Cycling Cancellation-Automaton for Sentence Well-Formedness*. International Computation Centre Bulletin **5**, 69–94.
- [Hudson1984] R. Hudson. 1984. *Word Grammar*. Blackwell.
- [Lambek2008] J. Lambek. 2008. *From Word to Sentence*. Polimetrica, Milan.
- [Landauer1997] T. Landauer, and S. Dumais. 2008. *A solution to Platos problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge*. Psychological review.
- [Manning2008] C. D. Manning, P. Raghavan, and H. Schütze. 2008. *Introduction to information retrieval*. Cambridge University Press.
- [Mitchell and Lapata2008] J. Mitchell and M. Lapata. 2008. *Vector-based models of semantic composition*. Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics, 236–244.
- [Montague1974] R. Montague. 1974. *English as a formal language*. Formal Philosophy, 189–223.
- [Nivre2003] J. Nivre 2003. *An efficient algorithm for projective dependency parsing*. Proceedings of the 8th International Workshop on Parsing Technologies (IWPT).
- [Saffran et al.1999] J. Saffron, E. Newport, R. Asling. 1999. *Word Segmentation: The role of distributional cues*. Journal of Memory and Language **35**, 606–621.
- [Schutze1998] H. Schuetze. 1998. *Automatic Word Sense Discrimination*. Computational Linguistics **24**, 97–123.
- [Smolensky1990] P. Smolensky. 1990. *Tensor product variable binding and the representation of symbolic structures in connectionist systems*. Computational Linguistics **46**, 1–2, 159–216.
- [Steedman2000] M. Steedman. 2000. *The Syntactic Process*. MIT Press.
- [Widdows2005] D. Widdows. 2005. *Geometry and Meaning*. University of Chicago Press.
- [Wittgenstein1953] L. Wittgenstein. 1953. *Philosophical Investigations*. Blackwell.